



AF
JPW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Douglas P. Brown et al.

§ Art Unit: 2166

Serial No.: 09/977,038

§

§

§

§

§

§

§

§

§

Filed: October 12, 2001

For: Index Selection in a Database
System

Examiner: Isaac M. Woo

Atty. Dkt. No.: 10150 (NCR.0063US)

Mail Stop Appeal Brief-Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

APPEAL BRIEF PURSUANT TO 37 C.F.R § 41.37

Sir:

The final rejection of claims 1-21, 40, and 43-46 is hereby appealed.

I. REAL PARTY IN INTEREST

The real party in interest is NCR Corporation.

II. RELATED APPEALS AND INTERFERENCES

None.

III. STATUS OF THE CLAIMS

Claims 1-21, 40, and 43-46 have been finally rejected and are the subject of this appeal.

Claims 22-39 and 41-42 have been cancelled.

Date of Deposit:

I hereby certify under 37 CFR 1.8(a) that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage on the date indicated above and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313.

Ginger Yount

IV. STATUS OF AMENDMENTS

The Amendment after final (mailed January 31, 2006) will be entered for purposes of appeal, as agreed in a telephonic interview between the undersigned and Examiner Hosain Alam on March 28, 2006.

As also agreed in the telephonic interview of March 28, 2006, the rejections under 35 U.S.C. § 101 of claims 1 and 7 were overcome, and the objection of claim 1 was overcome by the Amendment after final.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the independent claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element.

Independent claim 1 recites a test system (Fig. 1:10) comprising:

at least one processor (Spec., p. 47, lines 12-20);

an emulation module (Fig. 1:18) executable on the at least one processor to receive environment information of a database system separate from the test system, the emulation module to emulate an environment of the database system based on the environment information (Spec., p. 7, lines 1-11; p. 12, lines 10-13);

a first module (Fig. 1:50; Fig. 9:310) executable in the emulated environment and adapted to receive a set of queries and to provide a set of candidate indexes for the set of queries, the first module adapted to eliminate one or more candidate indexes based on one or more predetermined criteria (Spec., p. 7, lines 24-26; p. 13, line 20-p. 14, line 2; p. 16, lines 29-30; p. 20, lines 23-30; p. 22, lines 19-28; p. 24, line 11-p. 25, line 21); and

a second module (Fig. 9:319) executable in the emulated environment and adapted to generate a recommended index from the set of candidate indexes (Spec., p. 7, lines 24-26; p. 8, lines 15-30; p. 9, lines 17-26; p. 13, line 20-p. 14, line 2; p. 16, lines 29-30; p. 21, lines 14-18; p. 23, lines 14-15; p. 30, lines 26-31).

Independent claim 7 recites a system (Fig. 1:10, 14; Fig. 9:300) comprising:

at least one processor (Spec., p. 47, lines 12-20);

a first module (Fig. 1:50; Fig. 9:310) executable on the at least one processor to receive a set of queries and to provide a set of candidate indexes for the set of queries, the first module adapted to eliminate one or more candidate indexes based on one or more predetermined criteria (Spec., p. 7, lines 24-26; p. 13, line 20-p. 14, line 2; p. 16, lines 29-30; p. 20, lines 23-30; p. 22, lines 19-28; p. 24, line 11-p. 25, line 21); and

an optimizer (Fig. 1:18) adapted to generate a recommended index from the set of candidate indexes (Spec., p. 7, lines 24-26; p. 8, lines 15-30; p. 9, lines 17-26; p. 13, line 20-p. 14, line 2; p. 16, lines 29-30; p. 21, lines 14-18; p. 23, lines 14-15; p. 30, lines 26-31),

wherein the one or more predetermined criteria comprises a threshold change rate, the first module adapted to eliminate one or more candidate indexes having a change rate exceeding the threshold change rate (Spec., p. 13, lines 11-16; p. 20, lines 8-22; p. 24, line 13-p. 25, line 5).

Independent claim 40 recites an article comprising at least one storage medium containing instructions (Spec., p. 47, lines 12-29) that when executed cause a system (Fig. 1:10, 14; Fig. 9:300) to:

receive a set of queries (Spec., p. 9, lines 1-16);

generate a set of candidate indexes from the set of queries (Spec., p. 7, lines 24-26; p. 13, lines 25-27; p. 22, lines 19-20);

eliminate candidate indexes based on one or more predetermined criteria (Spec., p. 20, lines 23-30; p. 24, line 11-p. 25, line 21);

invoke an optimizer to perform cost analysis of the candidate indexes (Spec., p. 13, lines 22-25); and

use the cost analysis to select a recommended index for a database system (Spec., p. 7, lines 24-26; p. 8, lines 15-30),

wherein eliminating candidate indexes based on one or more predetermined criteria comprises at least one of:

eliminating candidate indexes that are changed with updates at a rate greater than a predetermined change rate threshold (Spec., p. 13, lines 11-16; p. 20, lines 8-22; p. 24, line 13-p. 25, line 5); and

eliminating a candidate index that is a subset of another candidate index (Spec., p 25, lines 6-17).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

As agreed to in the telephonic interview on March 28, 2006 between the undersigned and Examiner Alam, the Amendment after final has overcome the rejection under 35 U.S.C. § 101 and the claim objection under 37 C.F.R. § 1.75(c). Therefore, the only remaining ground of rejection to be reviewed on appeal is the following:

- A. Claims 1-21, 40, And 43-46 Rejected Under 35 U.S.C. § 102 Over U.S. Patent No. 6,366,903 (Agrawal).**

VII. ARGUMENT

The claims do not stand or fall together. Instead, Appellant presents separate arguments for various independent and dependent claims. Each of these arguments is separately argued below and presented with separate headings and sub-headings as required by 37 C.F.R. § 41.37(c)(1)(vii).

- A. Claims 1-21, 40, And 43-46 Rejected Under 35 U.S.C. § 102 Over U.S. Patent No. 6,366,903 (Agrawal).**

- 1. Claims 1-4, and 17-21.**

Independent claim 1 was rejected as being anticipated by Agrawal.

A claim is anticipated only if each and every element recited in the claim is found in a single prior art reference. M.P.E.P. § 2131 (8th ed., Rev. 3), at 2100-76. In this case, Agrawal clearly does not teach or suggest each and every element of claim 1.

Agrawal does not disclose a test system that has an emulation module to receive environment information of a database system *separate from the test system*, where the emulation module is to emulate an environment of the database system based on the environment information. As clearly recited in claim 1, this emulation module is part of the test system that is *separate* from the database system.

A clear error made by the Examiner in the rejection of claim 1 is that the Examiner never even made reference to the words “separate from the test system” recited in claim 1 (line 4) in either the final Office Action dated December 1, 2006, or the Advisory Action dated February 21, 2006. In the rejection of claim 1, the Examiner simply left out the words “separate from the test system” when paraphrasing claim 1. In fact, the Examiner has failed to provide any explanation of which elements of Agrawal are considered the database system and which elements of Agrawal are considered the *separate* test system. Ignoring express claim elements in making a rejection constitutes clear error.

Agrawal does not disclose an emulation module in a test system separate from a database system. The Examiner cited element 240 of Agrawal (depicted in Fig. 2 of Agrawal) as being the emulation module of claim 1. Note that element 240 of Agrawal is the configuration simulation and cost estimation module that is part of the optimizer of the database server 245 in Fig. 2 of Agrawal. Agrawal, 7:48-51. In other words, the module 240 is actually part of the database server (or database system) in Fig. 2 of Agrawal. Therefore, the module 240 cannot be

considered the emulation module of claim 1, which must be in a test system that is separate from the database system. In Agrawal, the module 240 is actually part of the database system.

Agrawal further discloses a candidate selection module 225 (also shown in Fig. 2 of Agrawal) that receives cost information regarding candidates relative to the workload and the server from the configuration and simulation cost estimation module 240. Agrawal, 7:51-53. This candidate selection module 225 is part of a group of components identified generally as 210 in Fig. 2 of Agrawal. However, there is no indication in Agrawal that any of the components indicated generally as 210 performs the following task: emulate an environment of the database system based on the received environment information of the database system. There is absolutely no indication anywhere within Agrawal that any emulation is being performed by the components in 210 in Fig. 2 of Agrawal.

It is therefore clear that Agrawal does not disclose the test system of claim 1, which test system is separate from the database system and includes the recited emulation module, first module, and second module.

In view of the foregoing, it is respectfully submitted that claim 1 is not anticipated by Agrawal. Dependent claims of claim 1 are allowable for at least the same reasons.

Therefore, reversal of the final rejection of the above claims is respectfully requested.

2. Claim 5.

Claim 5 depends indirectly from claim 1, and is thus allowable for at least the same reasons as claim 1. Moreover, claim 5 recites that statistics used by the second module of claim 1 are based on a scan of a sample of one or more tables, where the sample is less than all the rows of the one or more tables. The Examiner cited Fig. 2 and column 14, lines 40-54 of Agrawal as disclosing this feature. Fig. 2 of Agrawal does not show the scanning of a sample of

one or more tables, where the sample is less than all the rows of the one or more tables. The cited passage in column 14 of Agrawal refers to estimating the size of a materialized view, where the estimation is based on multiplying an estimated number of rows in the materialized view definition by the total width of all the columns in the materialized view. There is clearly no teaching in this cited passage of statistics based on a scan of a sample of one or more tables, where the sample is less than all the rows of the one or more tables.

For the foregoing reasons, reversal of the final rejection of the above claim is respectfully requested.

3. Claim 6.

Claim 6 depends from claim 5, and is thus allowable for at least the same reasons as claims 1 and 5. Moreover, claim 6 recites that the test system further comprises a graphical user interface to receive an indication of a user-specified size of the sample. The Examiner cited the same passages of Agrawal against claim 6 that the Examiner cited against claim 5. There is clearly no teaching of a user interface in Agrawal for receiving an indication of a user-specified size of the sample (of one or more tables for scanning to produce the statistics).

In view of the foregoing, the final rejection of the above claim should be reversed.

4. Claims 9-16.

Claim 9 depends from claim 1, and is thus allowable for at least the same reasons as claim 1. Moreover, claim 9 recites a second module that comprises an analysis module and an optimizer, where the analysis module is adapted to apply a genetic algorithm, and the analysis module is adapted to cooperate with the optimizer to generate the recommended index using the

genetic algorithm. Claims 10-16 are dependent from claim 9 and are thus allowable for at least this additional reason.

The Examiner cited Fig. 2, column 4, lines 42-52, and column 15, lines 36-59, of Agrawal as disclosing the analysis module to apply a genetic algorithm and to cooperate with the optimizer to generate the recommended index using the genetic algorithm. Fig. 2 of Agrawal shows no application of a genetic algorithm. The cited passage in column 4 of Agrawal does not disclose applying a genetic algorithm. The cited passage in column 15 of Agrawal also does not disclose application of a genetic algorithm. In fact, the term “genetic algorithm” is used nowhere in Agrawal.

Therefore, the final rejection of the above claims should be reversed.

5. Claim 46.

Claim 46 depends from claim 1, and is thus allowable for at least the same reasons as claim 1. Moreover, claim 46 recites that the environment information comprises cost-related information, statistics, and random samples from the database system. The Examiner cited Fig. 2, column 7, lines 47-67, and column 8, lines 1-61, of Agrawal as disclosing the environment information recited in claim 46. The cited passages in columns 7 and 8 of Agrawal refer to a simulation cost estimation module of the optimizer 240 that provides cost information regarding candidates relative to a workload. The cited passages also mention that the module 240 simulates the presence of indexes and materialized views that do not exist. However, there is absolutely no teaching in these passages of environment information that includes random samples from a database system.

Therefore, the final rejection of the above claim should be reversed.

6. Claims 7 and 8.

Independent claim 7 is also not anticipated by Agrawal. Claim 7 recites a first module to eliminate one or more candidate indexes based on one or more predetermined criteria, where the one or more predetermined criteria comprises a threshold change rate, and where the first module eliminates one or more candidate indexes having a change rate exceeding the threshold change rate.

In the rejection of claim 7, the Examiner cited column 10, lines 34-67, and column 15, lines 50-58, of Agrawal as disclosing the “threshold change rate.” 12/1/2005 Office Action at 7. The cited passage in column 10 refers to a predetermined cost threshold (10:63) that is used for pruning or eliminating “uninteresting” table-subsets, which are subsets of tables referenced in a query of a workload (8:45-46). The cost threshold identified in the column 10 passage is not a threshold change rate. Moreover, eliminating subsets of tables referenced by a query of a workload is *not* the same as eliminating one or more candidate *indexes* based on the threshold change rate.

The cited column 15 passage of Agrawal refers to a fraction f of the total storage to be allocated to indexes (15:44-45). The fraction of total storage for allocation to indexes is not a threshold change rate as recited in claim 7. Moreover, there is no teaching in column 15 of Agrawal that this fraction of total storage is used for eliminating candidate indexes.

In view of the foregoing, it is respectfully submitted that claim 7 is not anticipated by Agrawal. Dependent claims of claim 7 are allowable for at least the same reasons.

The final rejection of the above claims should be reversed.

7. Claim 40.

Independent claim 40 is also not anticipated by Agrawal, which does not disclose eliminating candidate indexes that are changed with updates at a rate greater than a predetermined change rate threshold. Also, there is no teaching Agrawal of eliminating a candidate index that is a subset of another candidate index.

The Examiner cited passages in column 10 and column 15 of Agrawal with respect to the element of eliminating candidate indexes that are changed with updates at a rate greater than a predetermined change rate threshold. As discussed above with respect to claim 7, these passages clearly do not disclose eliminating candidate indexes having a change rate greater than a change rate threshold.

Also, the Examiner cited the same passages of Agrawal against the element of eliminating a candidate index that is a subset of another candidate index in claim 40. Although the cited column 10 passage of Agrawal refers to criteria for eliminating uninteresting table subsets from the workload, there is no teaching in this passage of eliminating a candidate index that is a subset of another candidate index. Column 15 of Agrawal refers to determining the fraction f of total storage to be allocated to indexes – however, allocating a fraction of total storage to indexes is not the same as eliminating a candidate index that is a subset of another candidate index.

In view of the foregoing, reversal of the final rejection of the above claim is respectfully requested.

8. Claim 43.

Claim 43 depends from claim 40, and is thus allowable for at least the same reasons as claim 40. Moreover, claim 43 recites applying a genetic algorithm to select the recommended

index. The Examiner cited column 8, lines 10-30, of Agrawal as disclosing the applying of a genetic algorithm to select a recommended index. The cited passage refers to the simulation cost estimate module 240 simulating the presence of indexes and materialized views that do not exist. There is absolutely no mention in this passage of Agrawal, or anywhere else in Agrawal, of applying a genetic algorithm to select a recommended index.

In view of the foregoing, reversal of the final rejection of the above claim is respectfully requested.

9. Claim 44.

Claim 44 depends from claim 40, and is thus allowable for at least the same reasons as claim 40. Claim 44 further recites that the system that performs the acts recited in claim 40 is a test system that is separate from the database system, and that the instructions when executed cause the test system to perform the importing and emulating acts of claim 44. In the rejection of claim 44, the Examiner makes absolutely no mention of the words “test system separate from the database system.” Ignoring the word “separate” is clearly improper. It is clear that Agrawal does not disclose a test system that is separate from the database system, where the test system performs the acts recited in claims 40 and 44.

In view of the foregoing, reversal of the final rejection of the above claim is respectfully requested.

10. Claim 45.

Claim 45 depends from claim 44, and is thus allowable for at least the same reasons as claims 40 and 44. Moreover, claim 45 recites that the environment information comprises cost-related information, statistics, and random samples from the database system.

Consistent with Appellant's argument presented above with respect to claim 46, it is clear that the cited passages of Agrawal do not disclose environment information that includes random samples from the database system.

Therefore, reversal of the final rejection of the above claim is respectfully requested.

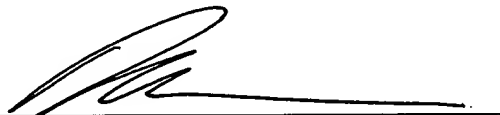
VIII. CONCLUSION

In view of the foregoing, reversal of all final rejections and allowance of all pending claims is respectfully requested.

Respectfully submitted,

Date: _____

June 27, 2006



Dan C. Hu
Registration No. 40,025
TROP, PRUNER & HU, P.C.
1616 South Voss Road, Suite 750
Houston, TX 77057-2631
Telephone: (713) 468-8880
Facsimile: (713) 468-8883

APPENDIX OF APPEALED CLAIMS

The claims on appeal are:

- 1 1. A test system comprising:
2 at least one processor;
3 an emulation module executable on the at least one processor to receive
4 environment information of a database system separate from the test system, the
5 emulation module to emulate an environment of the database system based on the
6 environment information;
7 a first module executable in the emulated environment and adapted to receive a
8 set of queries and to provide a set of candidate indexes for the set of queries, the first
9 module adapted to eliminate one or more candidate indexes based on one or more
10 predetermined criteria; and
11 a second module executable in the emulated environment and adapted to generate
12 a recommended index from the set of candidate indexes.
- 1 2. The test system of claim 1, wherein the set of queries comprises a set of SQL
2 statements.
- 1 3. The test system of claim 1, wherein the second module is adapted to generate at
2 least another recommended index from the set of candidate indexes.
- 1 4. The test system of claim 1, wherein the second module comprises an optimizer
2 that is adapted to use statistics.
- 1 5. The test system of claim 4, wherein the statistics are based on a scan of a sample
2 of one or more tables, the sample less than all the rows of the one or more tables.
- 1 6. The test system of claim 5, further comprising a graphical user interface to receive
2 an indication of a user-specified size of the sample.

1 7. A system comprising:
2 at least one processor;
3 a first module executable on the at least one processor to receive a set of queries
4 and to provide a set of candidate indexes for the set of queries, the first module adapted to
5 eliminate one or more candidate indexes based on one or more predetermined criteria;
6 and
7 an optimizer adapted to generate a recommended index from the set of candidate
8 indexes,
9 wherein the one or more predetermined criteria comprises a threshold change rate,
10 the first module adapted to eliminate one or more candidate indexes having a change rate
11 exceeding the threshold change rate.

1 8. The system of claim 7, wherein the first module is adapted to further eliminate a
2 candidate index that is a subset of another candidate index.

1 9. The test system of claim 1, wherein the second module comprises an analysis
2 module and an optimizer, the analysis module adapted to apply a genetic algorithm, the
3 analysis module adapted to cooperate with the optimizer to generate the recommended
4 index using the genetic algorithm.

1 10. The test system of claim 9, wherein the first module is adapted to provide the set
2 of candidate indexes by identifying the candidate indexes from the set of queries and
3 defining the set of queries in a database.

1 11. The test system of claim 10, wherein the analysis module is adapted to access the
2 database to retrieve the candidate indexes.

1 12. The test system of claim 10, further comprising a validation module adapted to
2 validate the recommended index in a database system.

1 13. The test system of claim 12, further comprising a user interface to receive user-
2 specified one or more indexes, the optimizer adapted to generate a cost associated with a
3 query plan based on the user-specified one or more indexes.

1 14. The test system of claim 13, wherein the user interface is adapted to receive a
2 user-specified percentage value, the system further comprising another module to collect
3 statistics based on a sample of rows of one or more tables, a size of the sample based on
4 the user-specified percentage value.

1 15. The test system of claim 14, further comprising another module adapted to
2 provide a hint on which table or tables statistics need to be collected.

1 16. The test system of claim 10, wherein the analysis module is adapted to access the
2 database to retrieve the candidate indexes.

1 17. The test system of claim 1, wherein the second module comprises an analysis
2 module and an optimizer, the analysis module adapted to apply a predetermined
3 algorithm, the analysis module adapted to cooperate with the optimizer to generate the
4 recommended index using the predetermined algorithm.

1 18. The test system of claim 17, wherein the analysis module is adapted to submit
2 candidate indexes to the optimizer, the optimizer adapted to determine the cost of one or
3 more of the queries based on the candidate indexes.

1 19. The test system of claim 18, wherein the optimizer is adapted to select the
2 candidate index associated with a lowest cost as the recommended index.

1 20. The test system of claim 1, wherein the set of queries comprises a workload
2 captured from the database system, and wherein the database system is a parallel system
3 having plural access modules, the environment information containing information
4 regarding the parallel system and plural access modules.

1 21. The test system of claim 20, wherein the optimizer is adapted to compute costs for
2 the candidate indexes in the emulated environment of the database system.

1 40. An article comprising at least one storage medium containing instructions that
2 when executed cause a system to:

3 receive a set of queries;

4 generate a set of candidate indexes from the set of queries;

5 eliminate candidate indexes based on one or more predetermined criteria;

6 invoke an optimizer to perform cost analysis of the candidate indexes; and

7 use the cost analysis to select a recommended index for a database system,

8 wherein eliminating candidate indexes based on one or more predetermined
9 criteria comprises at least one of:

10 eliminating candidate indexes that are changed with updates at a rate
11 greater than a predetermined change rate threshold; and

12 eliminating a candidate index that is a subset of another candidate index.

1 43. The article of claim 40, wherein the instructions when executed cause the system
2 to apply a genetic algorithm to select the recommended index.

1 44. The article of claim 40, wherein the system is a test system separate from the
2 database system, the instructions when executed causing the test system to:

3 import environment information regarding the database system;

4 emulate an environment of the database system based on the imported
5 environment information,

6 wherein the generating, eliminating, invoking, and using acts are performed in the
7 emulated environment.

1 45. The article of claim 44, wherein the environment information comprises cost-
2 related information, statistics, and random samples from the database system.

- 1 46. The article of claim 1, wherein the environment information comprises cost-
- 2 related information, statistics, and random samples from the database system.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.